

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 384 339 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention
of the grant of the patent:
02.04.1997 Bulletin 1997/14

(51) Int Cl.⁶: **G06F 9/46**

(21) Application number: **90103115.3**

(22) Date of filing: **19.02.1990**

(54) Broker for computer network server selection

Makler für die Auswahl von Rechnernetzwerkservern

Courtier pour la sélection de serveur de réseau d'ordinateur

(84) Designated Contracting States:
AT BE CH DE FR GB IT LI LU NL SE

(30) Priority: **24.02.1989 US 314853**

(43) Date of publication of application:
29.08.1990 Bulletin 1990/35

(73) Proprietor: **DIGITAL EQUIPMENT CORPORATION**
Maynard Massachusetts 01754-1418 (US)

(72) Inventors:

- **Pitkin, Richard P.**
Lowell, Massachusetts 01852 (US)
- **Morency, John P.**
Chelmsford, Massachusetts 01824 (US)

(74) Representative: **Betten & Resch**
Reichenbachstrasse 19
80469 München (DE)

(56) References cited:
EP-A- 0 248 403

EP 0 384 339 B1

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

Description

Field Of The Invention

This invention relates generally to the allocation of resources within a computer network architecture and, more particularly, to the use of a broker mechanism which responds to a client's or end user's request for some service and suggests a server to the client capable of supplying the requested service.

Background Of The Invention

Present day computer network architectures have become more and more complex. For example, a network may comprise multiple devices, such as several computer systems, terminals, peripherals, etc., all linked together in order to exchange information and share resources.

Other more complex networks may comprise thousands of devices of varying technical capabilities. In many applications, both local area networks (LANs) and wide area networks (WANs) are connected to form parts of a single network. For example, a network can include a local area component, a wide area component and a component involving communication between computers of different vendors, such as a multi-vendor communication network.

The process for designing computer networks of this type involves a complex analysis of the customer requirements, the price factor for the various products, the necessary communication lines, and other technical and business considerations. The allocation and planned use of the various resources to support the network needs is now part of the normal capacity planning cycle done by a network manager when setting up the network for a customer.

Because of the potential to form complex networks, a computer network client has a wide range of resources available to supply requested service. This advantage, however, also adds to the problem of networks and resource management. It is therefore desirable to obtain delivery of these services through network service names. A service offers clients capacity to use in accessing resources in the network.

In these computer networks, a server, made up of both hardware and software, may be used as an agent between a client and one or more resources to deliver the requested service to the client. The server's software program provides "actions" for a client using one or more resources. These actions, for example, may be transmission of packets over a communication line - a communication server; computations in a problem - a CPU server; information access in a data base - a data base server; network addressing information distributed throughout a network and associated with a requested service - a name server; or other similar and multiple combinations of these functions.

Because networks may include multiple servers generally delivering the same action, a client may access any one of these servers with equal delivery of the action. One of the disadvantages in the prior art is the inability to group these servers providing the same action. This grouping of servers can be seen as a "network service".

To overcome shortcomings of the prior art an ideal system for a client to request a service in a complex network should fulfill two requirements. First, the request must include generic accessing of a particular network service name to provide a higher level interface abstraction than the current practice where each client knows of all server possibilities. This has the advantage of substantially simplifying the decision process for the client in the case where multiple servers can deliver the requested service. The only requirement is the network service name rather than the name of each and every server offering the requested service. Second, the requested service must be guaranteed to exist for the client within the server for the duration of a connection.

To supply a service to a client, it is known to use a broker mechanism coupled to the network for scheduling a server to a requesting client. The broker receives a request for access to a service from a client and transmits to the client the name of a server which can deliver the service to the client. One known type of broker operates by assigning an entire server to a client irrespective of the capacity needed by the client.

A problem with the above broker method is the inefficient use of network resources. Prior brokers must continuously monitor all of the servers coupled to the network, increasing the brokers' overhead and the network traffic. Because the resource capacities of the servers and the necessary performance level to be allocated to the client are dynamic factors, a broker mechanism which can dynamically allocate servers to clients using a predetermined policy is needed.

European publication A30 248 403 teaches a complex computer system having several work stations and a plurality of host computers. Each work station has a built-in process to determine a server process required to implement a processing request. A communication control establishes a communication path between an associated server process and an applicable host. There is, however, no broker system in this prior art to allocate a plurality of servers.

SUMMARY OF THE INVENTION

The present invention is a broker method and apparatus which, on one hand monitors the dynamic status of a set of servers and, on the other hand, responds to requests from accessing clients concerning which member of that server set is capable of providing the requested service.

The invention, in its broad form, resides in a computer-implemented method for allocating a plurality of

servers in a network, and a device therefor, as recited in claims 1 and 10 respectively.

The service limitation for the requested service are preferably established as a matter of policy during the network design and modeling process by a system or network manager. Based upon the policy, the broker thus suggests to the client a server which is best able to satisfy the client's service request. The broker enables a client to use a service even though the client is unaware of the presence or absence of any individual servers within the network.

According to one aspect of the invention, the client then requests the service from the recommended server, and the server is responsible for granting the request only if the server currently has the required capacity available for that service. Since the server makes the final determination of whether it has the available capacity to provide the requested service, there is no risk of overloading a server because the broker has out of date status information - for example, as when a server becomes busy subsequent to its last status message transmitted to the broker.

As indicated above, the suggestion of a server preferably is based on the network policy developed by the network manager for the particular network. Because of the complex nature of the present networks, a modeling process is used to develop a local policy for each server in the network to deliver a service, based on the individual customer requirements for the network. The collection of local policies determines the overall network policy for a given service.

In most cases, the network policy is based on the servers' capacity to deliver a given service. The capacity of the given service, however, can be changed with the addition or subtraction of servers to the network. In either case, the changes are made known to the broker and are transparent to the clients in the network.

Upon startup, the broker develops a linked list data structure containing sets of server entries for each offered service corresponding to the local policy obtained for each server. According to one aspect of the invention, at any given time, the broker monitors a subset of those entries as a "preview window". The broker regularly rotates different servers into and out of the previous window. The broker suggests an entry in the preview window, upon receiving a client request, provided the entry has the available resources for the client.

An advantage of the preview window is to substantially decrease the number of servers which must be monitored by the broker in a network. Because the broker maintains current status information on only a subset of the servers during a given time interval, the number of status messages required to be transmitted over the network is reduced. This results in a reduction of computational overhead for the broker and communications overhead on the network.

In this embodiment, the broker assigns servers in a round robin fashion to ensure that the loss of a single

server does not have a major impact upon clients that request access at similar times. The application of the round robin server distribution, however, is regulated by use of a "scan weight" parameter for each server. The scan weight is defined as the number of client requests which can be satisfied by a server before that server is removed from the preview window.

The advantages of the broker can be achieved with only a minimal increase in connect time from client to server. This is because the broker uses the preview window to monitor server status information prior to receiving any client requests. The broker also uses the scan weight value to reduce the frequency by which the preview window is changed thus further increasing accurate server suggestions by the broker. Therefore, the incremental client connect time, as opposed to a direct client-server connection, is only increased by the time needed to contact the broker and receive a suggestion.

In a further embodiment, the broker suggests to a client both the top server entry and the next active server entry in the preview window. By suggesting two possible servers, the broker compensates for the case in which a failure of the first server is alleviated without the need to recontact the broker, thus decreasing the load on the network.

Another embodiment of the present invention makes use of multiple brokers for suggesting servers to multiple different clients. The ability to have multiple brokers running simultaneously provides a highly reliable service by alleviating any single point of failure.

Brief Description Of The Drawings

Figure 1 is a block diagram illustrating the environment of the present invention.

Figure 2 is a block diagram of a network architecture including the present invention.

Figure 2A is a block diagram conceptually illustrating the broker mechanism of the present invention.

Figure 3 is a flow chart describing the modelling of a network policy according to the present invention.

Figure 4 is a conceptual block diagram of the broker mechanism in Figure 2.

Figures 5 and 5A are examples showing the operation of server datastructures used in the present invention.

Figures 6 and 6A are flow charts describing the operation of the broker mechanism of the present invention.

Figure 7 is a conceptual block diagram of an embodiment of the present invention.

Detailed Description of the Invention

I. Network Environment

Referring to Figure 1, there is shown a model block diagram of a system enhanced by the present invention.

In a network 5, a plurality of users or clients, generally designated 10, each have access to a plurality of servers, generally designated 20. Some number of the servers 20 can provide access to services A_1-A_n requested by one of the clients 10. When the client 10 requests access to a service via one of the servers 20, a connection is made from the selected client 10 through the server 20. As can be seen in Figure 1, a plurality of servers 20 are available to make the requested connection. Therefore, it becomes a problem to know which client 10 should connect to which server 20. Further, the resource capacity of the server 20 must be known in order to make an appropriate connection. Because the server resources have various and different capabilities, i.e., small CPU size, large CPU size, used and unused capacities, different serial line speeds, etc., different servers 20 are more appropriate for a particular client request. Also, it is necessary to know the service level desired by the client 10.

Figure 2 is an example of a network architecture incorporating the present invention. The network 5 includes a plurality of servers 21-27 and clients 11-19 coupled to a communications medium 6, e.g., an Ethernet. A broker, conceptually shown as block 30, is also coupled to the communications medium 6. Further, a distributed repository 42 or 44 is illustratively coupled to the communication medium 6 for storing network policy. These repositories may be replications of each other distributed throughout the network. Each server can provide access to various services A_1-A_n for a requesting client 11-19. In operation, the broker 30 receives the local policies from the repository 42 or 44 for building data structures for each service supported by the servers, processes client requests for a service, and responds in accordance with network policy with one or more server suggestions to the client.

Figure 2A conceptually illustrates the architectural diagram of Figure 2. A broker mechanism 30 is used to suggest to clients 11-19 an appropriate server 21-26 for delivering the requested service (service A1 or A2). Further, individual servers 23, 24 and 25 can provide more than one service. While figure 2A only illustrates two services, it is apparent that the broker 30 can suggest servers to clients requesting any number of different services. When a particular client, shown as client 13, requests access to service A1, the client 13 places a request with the broker 30 on path 54. The broker 30 suggests an appropriate server from server set 21-24 to the client 13.

II. Modelling

A modelling process is used to efficiently determine the allocation of resources when designing a network. The modelling is accomplished by the network manager both prior to implementing the network and upon making subsequent changes to the network. In the modelling process factors are taken into account such as whether

data transmissions during the use of a particular network service are of a bulk or interactive (bursty) nature and their effect on resource capacities. Because bulk data transfers usually have a high transfer rate i.e., an entire floppy diskette of data is transferred and, on the other hand, interactive data transfers generally are of a short, bursty nature, it is generally possible in a given server to support more interactive data transfers than bulk data transfers.

Figure 3 is a flow chart depicting the modelling process that occurs in the development of the network policy to be implemented by the broker mechanism for each service.

The modelling begins by determining the service characteristics that are desired by the customer, i.e. the various possible actions, and the parameters of the server in the network, i.e. the resources of the server. Both the server parameters and required service characteristics are inputs to a modelling process such as is described in "Processor - sharing queuing model for time-shared systems with bulk arrivals", Kleinrock et al, Networks v.1 n.1 pp. 1-13 (1971) or "Models for Computer Networks", Kleinrock, IEEE Int. Conference on Communications, Conf. Rec, Boulder, Co, June 9-11, Session 21 pp. 9-16 (1969).

III. Network Policy

The model produces a prediction of the measured performance characteristics of each server based on the performance desired for each service offered by that service. This prediction is compared to the desired performance for each service to check if the server meets the performance limitations. If the prediction matches the desired performance, then the local policy for that particular server has been obtained. However, if the prediction does not match the desired performance, then the local policy parameters of the server and/or the service characteristics are varied before being modeled again. This process continues until the server prediction matches the desired performance thus developing the local policy for the server.

Next, the modelling process checks whether a local policy has been established for each server in the network. If not, then the above process is repeated. Once all local policies have been determined, the local policies are collected to form a network policy for the entire network. The network policy is stored in the distributed repository 42 or 44 as shown in Figure 2.

When customers set up their network, they therefore establish a network policy to allocate resources for each service. Policy decisions must be cognizant of all of the available factors made prior to implementing the system, i.e., based on the physical constraints of the hardware, the communication lines, the price factor for the product, the usage load, i.e., heavy, medium, light, fast CPUs, slow CPUs and other appropriate factors. A broker implements the network policy when suggesting

an appropriate server to maximize the efficiency of the complex network.

IV. Broker

An advantage of the invention allows the broker 30 (Figure 2A) to operate when there is: a network policy for a given service; a method to easily determine the current usage of the server relative to the established local policy; and servers which can enforce their local policy e.g. by rejecting any connection attempted beyond their local policy limit. These factors are used by the broker mechanism 30 to suggest a server 21-26 to a requesting client 13 based on the collection of local server policies in the broker.

The data structures of the broker mechanism 30 are formed from an application of the distributed repository storing the network policy for each service from which the broker linked list data structures can be built. Many conventional means such as name server computer programs are available to facilitate the creation of these broker linked list data structures. The distributed repository is a non-volatile storage area which holds the collection of local policies and the attributes of how each supporting server supports a service.

The repository is distributed throughout the network. As part of the broker startup procedures, the broker 30 extracts from the repository attribute information about each of the services it implements, i.e., service A_1 - A_n as shown in Figure 1. The repository is necessary to store the network policy in the event of a broker failure as the broker functions to compare current server capacity to network policy.

Figure 4 is a conceptual diagram showing an example of the data structures produced by the code in the broker mechanism 30. The data structure shown inside the broker 30 are linked lists of entries created by the broker code.

The broker obtains the entries to create the data structures from the distributed repository 42. A service list 71 is generated for those services found in the repository 42. Further, a server list 74 is created for each service in the service list 71 to hold the local policy information for the particular servers 21-26 that support the service. Lastly, the broker creates a "server status block" which contains a number of server connection entries 922, 923, 924 and 926, there are being one connection entry in the server status block for each server 22, 23, 24 and 26 which currently is in the "preview window" of at least one service. (Preview status will be defined below.) Each connection entry stores the current status information for its corresponding server. The policy information contained in the server list and service lists remains static, while the server status block information is of a dynamic or volatile nature.

A server connection section 75 of the broker 30 establishes communication paths 31-34 with the servers referred to in the preceding paragraph. The communi-

cation paths 31-34 allow the broker to poll each coupled server (22, 23, 24, 26) to receive its status. The status of the servers 22, 23, 24 and 26 is stored in connection entries 922, 923, 924 and 926, respectively, within the server status block. In response to subsequent client requests for service, the broker examines these connection entries to determine each coupled server's capacity or availability to deliver a requested service, as will be described below with reference to Figures 5 and 5A. The server connection code 75 therefore supports a list of connection entries 922-926 that map (80-84) to a select number of server entries in the server lists 74 for each of the services in service list 71. As shown in Figures 4, a server 24 can overlap between service offerings. The server connection code 75 may therefore have several server lists 74 with entries pointing to a single server connection. This pooling of server connections is important for providing multiple service offerings from one broker location.

1. Preview Window

During the broker's startup, data structures are developed in the broker. The data structures include a service list 71 and server lists 74 having sets of server entries (121-125) and (223-226) corresponding to the servers (21-25) and (23-26) for each service (A_1 & A_2) in the service list 71. The data structure information is provided from the repository 42. The broker 30 then monitors, via couplings 80-84, subsets (22, 23 & 24) and (24 & 26) of the servers from the above sets to form a preview window for each service. The size of the preview window is predetermined according to the network policy and stored as an attribute of the network policy in the repository 42. The preview windows 72 ensure that a sufficient number of servers, i.e., based upon the network policy for each service, are coupled and providing status to the broker 30 for each service before a client request is made. The size of the preview window generally reflects the number of servers necessary to satisfy the average rate of client requests received by the broker. By using preview windows, the connection time from client to servers is reduced by accumulating status and policy information in the broker in advance of the client request.

Without using the preview windows 72, the preconnection time for a client to receive data from a service would be the accumulation of: (client connect time to broker) + (broker connect time to server location) + (server startup time) + (accumulation of connect times needed to find a server location having unassigned capacity). By monitoring the resources of the servers contained in the preview windows via communication paths 31-34 associated with each service, the broker 30 can suggest the top entry in the preview window to a client requesting that particular service. The suggestion is made provided the entry has the available resources. If an entry does not have the available resources it is re-

moved from the preview window. Servers are thus regularly rotated in and out of the preview window.

Figures 5 and 5A illustrate by example the data structures within each server (shown here as servers 1 and 2) to indicate server capacity and status to deliver a requested service. The status information is provided to the broker 30 for comparing the capacity to the local policy for each server as shown by communication paths 31-34 in Figure 4.

In Figure 5, server 1 is shown supporting two services 101 and 102 having service names A_1 and A_2 , respectively. Similarly, server 2 supports services 103 and 104 having service names A_1 and A_2 . Associated with each service 101-104 is a client capacity number 111-114 indicating the number of clients that can be processed by each server for the particular service. The client capacity number reflects the server's local policy for each service. As shown in Figure 5, the client capacity number can be a scalar value which is reduced as the server increases its load. The current value of the capacity number 111-114 is therefore checked by the broker whenever the broker tries to use the server.

A second implementation of the database of server 1 is shown in Figure 5A wherein client slots 105-108 are assigned to each service A_1 and A_2 (101 and 102). For example, assuming each server has four client slots, two client slots (105, 106 and 107, 108) can be allotted for each service 101 and 102, respectively in server 1. This is equivalent to a client capacity number of two for each service 101 or 102. Client slots allow the broker to determine the availability of servers by checking a flag indicating whether particular client slots are used.

However, a more powerful use of client slots is achieved in server 2 as shown in Figure 5A by allowing certain client slots to provide multiple services. Again, only four client slots are allocated for server 2. The first client slot 115 is dedicated for use with service A_1 (103) and the second client slot 118 is provided for service A_2 (104). In this case, however, the third and fourth client slots are made available to each service 116, 117, 119, 121. Therefore, if three requests are received for service A_2 and only one request for service A_1 , then by making the client slots available for use with each service, all of the requests can be satisfied. Again, this is implemented by the server connection section 75 (Figure 4) checking flags to determine the availability of each client slot. These flags indicate the status for the servers against which the local policies are checked in connection entries 922-926.

Through the use of the preview window and the overlapping of servers to provide different services, the broker mechanism 30 thus reduces the number of active network connections (31-34) necessary to poll the servers to obtain their status. Otherwise, active connections would be required from all of the servers in the network. Therefore, only a reasonably managed subset of servers is maintained by the broker 30 as active connections to reduce message traffic on the network and to simplify

the computational overhead in the broker.

B. Scan Weights

Referring again to Figure 4, because the capacity of each server to deliver a given service may not be equal depending upon the network policy, an additional parameter termed scan weight 73 is added. The scan weight 73 is the number of client requests which can be satisfied from a server before the particular server is removed from the preview window 72. Each server is assigned a particular scan weight value for each service by the network manager. The scan weight 73 allows the network manager to apply the network policy to more efficiently assign clients 10 to servers in a round-robin manner. The values chosen for the scan weight are developed as part of the capacity planning modelled in Figure 3.

An example of scan weight operation is given in the server list 74 for service A_1 as shown in Figure 4. Five servers 21, 22, 23, 24, and 25 having server entries 121-125, respectively, are determined to have the following capacities to access a service for a client: server 21, 2 clients(10); server 22, 6 clients(10); server 23, 2 clients(10); server 24, 4 clients(10); and server 25, 2 clients(10). Without scan weight, the broker mechanism 30, using a purely round robin method of assigning servers, would reduce the list of five available servers 21-25 to just two available servers 22 and 24 after only 10 client requests have been suggested. This occurs because the servers having the least capacity, i.e. 21, 23 and 25, are suggested as often as the other servers, thus quickly using their capacity. After the first ten requests, only servers 22 and 24 have any capacity remaining. Further, only one server 22 would be available after the next four client requests. This rapid reduction in the available servers can produce single failure points which may disrupt a network.

A scan weight parameter 73 is therefore determined based on an equal distribution of the server capacities to handle client requests. The scan weight 73 controls the number of clients assigned to a particular server when it is at the top of the preview window. In our example, an appropriate allocation of clients per scan weight may be assigned as follows: 21, 1 client(10), 22, 3 clients(10); 23, 1 client(10); 24, 2 clients(10); and 25, 1 (10) client. By using these scan weight values, the servers having the greatest capacity, i.e. 22 and 24, will be assigned to multiple clients before being removed from the preview window. Thus, the client requests are evenly distributed across all available servers 21-25 based on their capacity. In this example, four servers would still be available after the first ten client requests have been satisfied.

C. CLIENT CONNECTION and MANAGEMENT OF BROKER

The client connection section 70 of the broker 30 is the interface for providing server suggestions to the clients. A client's request for a service is received on path 54. Client connection section 70 checks to insure that the service is located in the service list 71 of supported services in the broker 30. Further, client connection section 70 performs the necessary steps to suggest a server to the accessing client.

The management section 76 of the broker 30 provides control and status information relative to the operation of the broker 30. For example, the management section 76 can enable and disable service names, examines the size of the preview window 72, displays information on the current entries in the preview window, or other management functions.

V. BROKER OPERATION

The operation of Figure 4 is described by the flow charts of Figures 6 and 6A. Figure 6 describes the inclusion of the local policies into the operating broker mechanism 30. Figure 6A describes the operation of the broker upon receiving a client request for a server.

In operation upon broker 30 start-up, the services i.e., service A_1 , A_2 or A_3 , are obtained from the network list in the distributed repository 42 functioning as a directory for all of the services. Once the service is found, the broker 30 begins to build a service list data structure 71 which includes the attributes for each of the services. Next, the server information for the particular service is obtained from the repository 42. If no servers are found, then a new service is obtained from the repository 42 and the process is repeated. However, if a server is found, then a server list data structure 74 is built for the particular service. The server list data structure 74 contains the local policy for each server supporting the service and the scan weight 73 for each server in the server list 74. Once all of the services stored in the repository are found, then the server connection section 75 begins to build the status connections used for the preview windows 72 for each service.

Referring to Figures 4 and 6A, the operation of the broker upon receiving a client request for a service is described. The broker 30 receives a message from the client 13 via path 54 containing the name of the service requested, i.e., service A_1 , A_2 or A_3 . The client connection section 70 receives the request and checks whether the service list 71 exists in the broker. If the service list 71 exists, the broker 30 checks whether the particular service requested (in this example service A_2) is contained in the service list 71. The broker 30 then obtains the top entry 226 in the preview window 72 for service A_2 . Next, the broker checks whether the connection entry 926 is active, i.e., is there a connection from broker to server as on path 34. If it is not active, then the entry

226 is removed from the preview window 72. However, if the entry 226 is active, then the local policy information for the associated server 26 is obtained from entry 226 in the server list 74.

The local policy limit is then checked in connection entry 926 against the status information received for the server 26 on path 34. This is accomplished as described above in figures 5 and 5A by checking either the client capacity number or the availability of client slots, depending upon the type of status indication that is used.

The availability of server 26 to support service A_2 is checked to make sure that its policy limit is not exceeded. If the server 26 has the necessary capacity as determined by either the client slots or the client capacity number, i.e. the current value is less than the client capacity number or there are available client slots as indicated by the flags, then, a message is sent to the client 13 suggesting server 26 to be used to deliver service A_2 . The scan weight value 73 for server entry 226 is decreased and is checked to determine if the entry 226 should be removed from the preview window 72. The broker continues by refilling the preview window if necessary.

If, however, server 26 does not have the available capacity, then the next entry 224 in the preview window 72 is checked. This process continues until either an entry is found or no server is determined to be available. It is thus apparent that broker 30 need only check the status of a server when a request for a particular service is received and that particular server is in the preview window. For each entry in the preview window not having the necessary capacity or having been previously assigned, the entry is rotated out of the preview window and another entry refills the preview window. It may occur, however, that no servers are available to refill the preview window, in which case the size of the preview window would be reduced.

Under normal operations, the broker mechanism efficiently operates to receive client requests and suggest servers. However, if at any time during a request, the service list, name of the service or preview window do not exist, then the broker 30 will send an error message to the client 12.

An alternate embodiment of the present invention has the broker 30 suggest to the client 12, not only the top server entry 226 but also the next active server entry 225 in the preview window 72.

In operation, upon receiving a client request, the broker again would check the entry in the preview window. Once the broker finds an entry having the available capacity, it suggests that entry along with the next entry in the preview window. In this embodiment, only the scan weight of the first entry is decreased. The client then will attempt to access the service through the first entry. If however, the first entry fails for any reason, then the client attempts to use the second server entry without having to reconnect to the broker.

When using this embodiment, the broker does not

need to continuously check the status of the two entries. The status will be updated the next time the broker checks on those particular server entries.

The use of two entry suggestions is useful in the event of a failure by the first server to deliver the requested service. Failure of a server can occur due to other clients accessing the server directly thus using its capacity, failure of the hardware, line damage, etc. Thus, suggesting two servers provides for the situation in which the first server failed during the client's connection time to the server using the broker. The inclusion of a second server suggestion provides the client with an alternate server without having to recontact the broker. Thus, client connection time can be decreased by using two server suggestions.

As shown conceptually in Figure 7, another embodiment of the present invention incorporates multiple brokers 30 and 31 for suggesting servers to different clients 11-19. The brokers 30-31 can be replications of each other supporting the same services. The brokers operate independently of each other to provide a high availability factor to the clients. An attendant tradeoff for the availability is an increase in overhead on the network as status information for the monitored servers in the preview window must go to each broker 30, 31. If however, the rate of client requests to the brokers 30, 31 is slow, then the brokers can adjust the rate at which they poll the servers for their status in the preview window. By reducing the polling rate in proportion to the number of brokers, the message traffic on the network will remain substantially the same as for a single broker while still providing high availability and eliminating the broker as a single point of failure.

Using multiple brokers, the scan weight value increases is a linear function of the number of brokers in the network provided the network policy is consistent among the brokers. This occurs because each broker maintains a separate preview window and yet the enforcement of the local policies is carried out by each individual server. Because servers for the same service will be contained in each broker's preview window and the brokers are independent of each other, the servers will be rotated independently of the other brokers preview window. Thus, because the server has the same scan weight in two separate preview windows, the scan weight value is effectively doubled.

The ability to have multiple brokers 30, 31 operating simultaneously provides a highly reliable system by alleviating any single point of failure. Further, the clients gain an increased opportunity to use a broker to suggest a server for delivering the requested service.

A further embodiment of the present invention uses multiple server lists having different priorities. The different lists each have their own preview windows. The use of prioritized server lists allows the client to obtain service when the servers in the higher priority service list are unavailable. This is done by automatically redirecting the client requests to the lower server list to pro-

vide the service, albeit at a less desirable service level, without the client needing to make a request for access to an additional service name.

Claims

1. A computer-implemented method for allocating a plurality of servers (21-27) in a network (5), each server having an available resource capacity, to a plurality of clients (11-19) for delivering a plurality of services (A_1 - A_n) to said clients, the method comprising the steps of:

- a) developing a network policy for said plurality of servers (21-27) by collecting a local policy for each of said servers, wherein the local policy for each server is based on an overall capacity of each server to deliver each of the plurality of services (A_1 - A_n) to said clients (11-19);
- b) receiving client requests for said services in at least one broker (30, 31);
- c) making a server suggestion, from the broker (30, 31) to one of said clients (11-19) making a request for one of said services (A_1 - A_n), the server suggestion identifying a suggested one of said servers based on the network policy and available resource capacities of said servers, said suggested one of said servers having the available resource capacity to deliver said one of said services;
- d) making a request for said one of said services (A_1 - A_n) from said one of said clients (11-19) to said suggested one of said servers (21-27) in response to said one of said clients receiving the server suggestion of step (c) from the broker (30, 31);
- e) operating the suggested one of said servers in accordance with its respective local policy to reject the request for said one of said services (A_1 - A_n) of step (d) when the request exceeds a local policy limit for the respective local policy of the suggested one of said servers; and
- f) operating the suggested one of said servers (21-29) in accordance with its respective local policy to accept the request for said one of said services (A_1 - A_n) of step (d) when the request does not exceed the local policy limit for the respective local policy of the suggested one of said servers.

2. A method according to claim 1 wherein step (c) further includes:

- a) creating a service list (71), within said broker (30, 31), of available services from said network policy;
- b) creating a respective server list (74) for each

- service in the service list, each server list containing a server entry for each server having the overall capacity to deliver a respective service; and
- c) monitoring, for each service in the service list (71), a subset of said server entire in said server list for suggesting said servers represented in said subset to said clients.
3. A method according to claim 2 wherein said step of monitoring said subset further includes the steps of:
- a) establishing a data path (31-34), the data path providing connectivity between each server represented in the subset and the broker, the data path for transmitting information regarding the available resource capacity of each server represented in the subset to said broker; and
- b) determining whether a first server of the servers represented in the subset has an available client slot (105-108) to deliver the service requested by said clients based upon the available resource capacity of said first server.
4. A method according to claim 3 wherein said step of making a suggestion is made when one of said servers being monitored has an available client slot (105-108) to deliver said service.
5. A method according to claim 4 further comprising the steps of:
- a) disconnecting the data path (31-34) from said first server when said data path indicates no available client slots (105-108); and
- b) subsequently establishing with the broker (30, 31) a data path from another of said plurality of servers in said server list (74) supporting said service having available client slots.
6. A method according to claim 5 further comprising the steps of:
- a) developing a scan weight value (73) for each of said plurality of servers (21-27) supporting each service based on said network policy;
- b) creating a scan weight entry containing the scan weight value for each server entry in said server list (74);
- c) decreasing the scan weight value (73) for a selected one of said servers represented in the subset after said selected one is suggested to said client; and
- d) removing said selected one from the subset when said scan weight value is zero.
7. A method according to claim 1, wherein said step of developing a network policy includes:
- a) determining a set of service characters for each of said services;
- b) determining a service capability for one of said servers as a function of said sets of service characteristics and a set of server parameters of said one of said servers;
- c) comparing the service capability with a desired performance for said one server to determine if there is a match;
- d) varying the sets of service characteristics and the set of server parameters until said comparing step produces a match;
- e) storing said set of server parameters and said set of service characteristics as the local policy for said one server upon determination of a match;
- f) repeating steps (a)-(e) for each of said plurality of servers; and
- g) collecting said local policies to obtain a network policy.
8. A method according to claim wherein said computer network (5) uses multiple brokers (30, 31).
9. A method according to claim 8 wherein said multiple brokers (30, 31) are exact replicas of each other and operating independently.
10. A device for allocating a plurality of servers (21-27), each having an available resource capacity, to a plurality of clients (11-19) for delivering one of a plurality of services (A_1 - A_n) to said clients, said servers and said clients being arranged in a computer network (5), the device suggesting a suggested one of the plurality of servers (21-27) to a requesting one of the plurality of clients (11-19), the requesting one of the plurality of the clients (11-19) making a request for one of the plurality of services (A_1 - A_n) to the suggested one of the plurality of servers, the suggested one of the plurality of servers operating to reject the request for the one of the plurality of services when the request exceeds a local policy limit for a respective local policy of the suggested one of the plurality of servers (21-27), the suggested one of the plurality of servers operating to accept the request for said one of the plurality of services (A_1 - A_n) when the request does not exceed the local policy limit for the respective local policy of the suggested one of the plurality of servers, the device comprising:
- a) at least one broker (30, 31) including:
- (i) means for receiving client request for said service; and
- (ii) means for suggesting by the broker (30, 31) one of said servers to one of said clients making a request based on an available resource capacity, said one server being suggested having

the available resource capacity to deliver said service.

11. A device according to claim 10, wherein said service means for suggesting further comprises:

1) means for creating a service list (71), within said broker (30, 31), of available services;
2) means for creating a server list (74), containing server entries of available servers for supporting each of said services in said server list; and
3) means for monitoring a subset of said server entries in said server list (74) for suggesting said servers represented in said server list to said clients (11-19).

12. A device according to claim 10 or 11 wherein said means for monitoring further comprises:

a) a data path (31-34) establishing connectivity from each server in a set of servers supporting one of said services to said broker (30, 31), the data path for indicating available client slots (105-108) for the server to said broker; and
b) means for determining whether said first server has an available client slot to deliver the service requested by the clients.

13. A device according to claim 12 wherein said means for suggesting sends a message to said client when one of said servers has an available client slot to deliver said service.

14. A device according to claim 13 further comprising:

a) a scan weight value (73) developed for each of said plurality of servers (21-27);
b) a scan weight entry containing the scan weight value stored in said broker (30, 31) for each server entry in said server list;
c) means for decreasing the scan weight value (73) for said first server in the preview window (72) after said first server is suggested to said client; and
d) means for removing said first server from the preview window (72) when said scan weight value is zero.

15. A device according to claim 12 wherein the device further comprises: means for developing a network policy including,

a) means for determining a service capability for each server (21-27) based on a plurality of sets of service characteristics and a set of local policy parameters;
b) means for comparing, for each of said serv-

ers, the service capability with a desired performance for each server to determine if there is a match;

c) means for varying, for each of said servers (21-27), the sets of service characteristics and the set of local policy parameters until said comparing step produces a match; and
d) means for storing, for each of said servers (21-27), said set of local policy parameters and said sets of service characteristics as the local policy for each server upon determination of a match;

means for suggesting by the broker (30, 31) one of said servers to one of said clients based on the network policy and the available resource capacity of said one of said servers.

20 Patentansprüche

1. Verfahren, das in einem Computer implementiert ist, zum Zuweisen mehrerer Server (21-27) in einem Netz (5), wovon jeder eine Kapazität für verfügbare Betriebsmittel besitzt, an mehrere Clients (11-19), um mehrere Dienste (A_1 - A_n) an die Clients zu liefern, wobei das Verfahren die folgenden Schritte enthält:

a) Entwickeln einer Netzpolitik für die mehreren Server (21-27) durch Sammeln einer Lokalpolitik für jeden der Server, wobei die Lokalpolitik für jeden Server auf einer Gesamtkapazität jedes Servers zum Liefern jedes der Dienste (A_1 - A_n) an die Clients (11-19) basiert;
b) Empfangen von Client-Anforderungen der Dienste in wenigstens einem Makler (30, 31);
c) Erzeugen eines Server-Vorschlags vom Makler (30, 31) an einen der Clients (11-19), der eine Anforderung eines der Dienste (A_1 - A_n) erzeugt, wobei der Server-Vorschlag einen vorgeschlagenen der Server auf der Grundlage der Netzpolitik und der Kapazitäten für verfügbare Betriebsmittel der Server identifiziert, wobei der vorgeschlagene der Server die Kapazität für verfügbare Betriebsmittel zum Liefern des einen der Dienste besitzt;
d) Erzeugen einer Anforderung des einen der Dienste (A_1 - A_n) von dem einen der Clients (11-19) an den vorgeschlagenen der Server (21-27) als Antwort auf den einen der Clients, der den Server-Vorschlag des Schrittes (c) vom Makler (30, 31) empfängt;
e) Betreiben des vorgeschlagenen der Server in Übereinstimmung mit seiner jeweiligen Lokalpolitik, um die Anforderung des einen der Dienste (A_1 - A_n) des Schrittes (d) zurückzuweisen, wenn die Anforderung eine Lokalpolitik-

- Grenze der entsprechenden Lokalpolitik des vorgeschlagenen der Server überschreitet; und
f) Betreiben des vorgeschlagenen der Server (21-29) in Übereinstimmung mit seiner jeweiligen Lokalpolitik, um die Anforderung des einen der Dienste (A_1-A_n) des Schrittes (d) anzunehmen, wenn die Anforderung die Lokalpolitik-Grenze für die entsprechende Lokalpolitik des vorgeschlagenen der Server nicht überschreitet.
2. Verfahren nach Anspruch 1, in dem der Schritt (c) ferner enthält:
- a) Erzeugen einer Dienste-Liste (71) im Makler (30, 31) der von der Netzpolitik zur Verfügung gestellten Dienste;
 - b) Erzeugen einer entsprechenden Server-Liste (74) für jeden Dienst in der Dienste-Liste, wobei jede Server-Liste einen Server-Eintrag jedes Servers enthält, der die Gesamtkapazität besitzt, um einen entsprechenden Dienst zu liefern; und
 - c) Überwachen einer Untermenge der Server-Einträge in der Server-Liste für jeden Dienst in der Dienste-Liste (71), um den Clients die Server vorzuschlagen, die in der Untermenge dargestellt sind.
3. Verfahren nach Anspruch 2, in dem der Schritt des Überwachens der Untermenge ferner die folgenden Schritte enthält:
- a) Erstellen eines Datenpfades (31-34), wobei der Datenpfad eine Verbindung zwischen jedem in der Untermenge dargestellten Server und dem Makler schafft und der Übertragung von Informationen bezüglich der Kapazität für verfügbare Betriebsmittel jedes in der Untermenge dargestellten Servers an den Makler dient; und
 - b) Bestimmen auf der Grundlage der Kapazität für verfügbare Betriebsmittel eines ersten Servers der in der Untermenge dargestellten Server, ob der erste Server einen verfügbaren Client-Schlitz (105-108) zum Liefern des von den Clients angeforderten Diensts besitzt.
4. Verfahren nach Anspruch 3, in dem der Schritt des Erzeugens eines Vorschlags ausgeführt wird, wenn einer der überwachten Server einen verfügbaren Client-Schlitz (105-108) zum Liefern des Diensts besitzt.
5. Verfahren nach Anspruch 4, ferner mit den Schritten:
- a) Trennen des Datenpfades (31-34) vom ersten Server, wenn der Datenpfad keine verfügbaren Client-Schlitz (105-108) angibt; und
 - b) anschließendes Herstellen eines Datenpfades von einem weiteren der mehreren Server in der Server-Liste (74), der den Dienst unterstützt und verfügbare Client-Schlitz besitzt, zum Makler (30, 31).
6. Verfahren nach Anspruch 5, ferner mit den Schritten:
- a) Entwickeln eines Abfragegewichtswerts (73) für jeden der mehreren Server (21-27), die jeden Dienst unterstützen, auf der Grundlage der Netzpolitik;
 - b) Erzeugen eines Abfragegewichtseintrags, der den Abfragegewichtswert für jeden Server-Eintrag in der Server-Liste (74) enthält;
 - c) Erniedrigen des Abfragegewichtswerts (73) für einen ausgewählten der Server, die in der Untermenge dargestellt sind, nachdem der ausgewählte Server für den Client vorgeschlagen worden ist; und
 - d) Entfernen des ausgewählten Servers aus der Untermenge, wenn der Abfragegewichtswert null ist.
7. Verfahren nach Anspruch 1, in dem der Schritt des Entwickelns einer Netzpolitik enthält:
- a) Bestimmen einer Menge von Diensteigenschaften für jeden der Dienste;
 - b) Bestimmen eines Dienst-Vermögens für einen der Server in Abhängigkeit von den Mengen von Dienst-Eigenschaften und einer Menge von Server-Parametern des einen der Server;
 - c) Vergleichen des Dienst-Vermögens mit einer gewünschten Leistungsfähigkeit des einen Servers, um festzustellen, ob eine Übereinstimmung vorhanden ist;
 - d) Verändern der Mengen von Dienst-Eigenschaften und der Menge von Dienst-Parametern, bis der Vergleichsschritt eine Übereinstimmung ergibt;
 - e) Speichern der Menge von Dienst-Parametern und der Menge von Dienst-Eigenschaften als die Lokalpolitik für den einen Server, wenn eine Übereinstimmung festgestellt wird;
 - f) Wiederholen der Schritte (a)-(e) für jeden der mehreren Server; und
 - g) Sammeln der Lokalpolitiken, um eine Netzpolitik zu erhalten.
8. Verfahren nach Anspruch 1, in dem das Computernetz (5) mehrere Makler (30, 31) verwendet.
9. Verfahren nach Anspruch 8, in dem die mehreren

Makler (30, 31) exakte Kopien voneinander sind und unabhängig arbeiten.

10. Vorrichtung zum Zuweisen mehrerer Server (21-27), wovon jeder eine Kapazität für verfügbare Betriebsmittel besitzt, an mehrere Clients (11-19) zum Liefern eines von mehreren Diensten (A_1-A_n) an die Clients, wobei die Server und die Clients in einem Computernetz (5) angeordnet sind, wobei die Vorrichtung einem anfordernden der mehreren Clients (11-19) einen vorgeschlagenen der mehreren Server (21-27) vorschlägt, wobei der anfordernde der mehreren Clients (11-19) einen der mehreren Dienste (A_1-A_n) vom vorgeschlagenen der mehreren Server anfordert, wobei der vorgeschlagene der mehreren Server in der Weise arbeitet, daß er die Anforderung des einen der mehreren Dienste zurückweist, wenn die Anforderung eine Lokalpolitik-Grenze für eine entsprechende Lokalpolitik des vorgeschlagenen der mehreren Server (21-27) überschreitet, und wobei der vorgeschlagene der mehreren Server in der Weise arbeitet, daß der die Anforderung des einen der mehreren Dienste (A_1-A_n) annimmt, wenn die Anforderung die Lokalpolitik-Grenze für die entsprechende Lokalpolitik des vorgeschlagenen der mehreren Server nicht überschreitet, wobei die Vorrichtung enthält:

a) wenigstens einen Makler (30, 31), der enthält:

- (i) eine Einrichtung zum Empfangen einer Client-Anforderung des Diensts; und
- (ii) eine Einrichtung, die auf der Grundlage einer Kapazität für verfügbare Betriebsmittel durch den Makler (30, 31) einen der Server für einen der Clients, die eine Anforderung erzeugen, vorschlägt, wobei der eine Server, der vorgeschlagen wird, die Kapazität für verfügbare Betriebsmittel besitzt, um den Dienst zu liefern.

11. Vorrichtung nach Anspruch 10, in der die Dienst Einrichtung zum Vorschlagen ferner enthält:

- 1) eine Einrichtung zum Erzeugen einer Dienst-Liste (71) der verfügbaren Dienste im Makler (30, 31);
- 2) eine Einrichtung zum Erzeugen einer Server-Liste (74), die Server-Einträge der verfügbaren Server enthält, die jeden der Dienste in der Server-Liste unterstützen; und
- 3) eine Einrichtung zum Überwachen einer Untermenge der Server-Einträge in der Server-Liste (74), um den Clients (11-19) die in der Server-Liste dargestellten Server vorzuschlagen.

12. Vorrichtung nach Anspruch 10 oder 11, in der die

Einrichtung zum Überwachen ferner enthält:

- a) einen Datenpfad (31-34) zum Herstellen einer Verbindung von jedem Server in einer Menge von einen der Dienste unterstützenden Servern zum Makler (30, 31), wobei der Datenpfad dem Makler verfügbare Client-Schlitz (105-108) des jeweiligen Servers anzeigt; und
- b) eine Einrichtung, die feststellt, ob der erste Server einen verfügbaren Client-Schlitz besitzt, um den von den Clients angeforderten Dienst zu liefern.

13. Vorrichtung nach Anspruch 12, in der die Einrichtung zum Vorschlagen an den Client eine Nachricht sendet, wenn einer der Server einen verfügbaren Client-Schlitz besitzt, um den Dienst zu liefern.

14. Vorrichtung nach Anspruch 13, ferner mit:

- a) einem Abfragegewichtswert (73), der für jeden der mehreren Server (21-27) entwickelt wird;
- b) einem Abfragegewichtseintrag, der den im Makler (30, 31) gespeicherten Abfragegewichtswert für jeden Server-Eintrag in der Server-Liste enthält;
- c) einer Einrichtung zum Erniedrigen des Abfragegewichtswerts (73) des ersten Servers in einem Vorausschaufenster (72), nachdem der erste Server dem Client vorgeschlagen worden ist; und
- d) einer Einrichtung zum Entfernen des ersten Servers aus dem Vorausschaufenster (72), wenn der Abfragegewichtswert Null ist.

15. Vorrichtung nach Anspruch 12, wobei die Vorrichtung ferner enthält:

eine Einrichtung zum Entwickeln einer Netzpolitik, mit

- a) einer Einrichtung zum Bestimmen eines Dienst-Vermögens für jeden Server (21-27) auf der Grundlage mehrerer Mengen von Dienst-Eigenschaften und einer Menge von Lokalpolitik-Parametern;
- b) einer Einrichtung zum Vergleichen des Dienst-Vermögens mit einer für jeden Server gewünschten Leistungsfähigkeit für jeden Server, um festzustellen, wenn eine Übereinstimmung vorhanden ist;
- c) einer Einrichtung zum Verändern der Mengen von Dienst-Eigenschaften und der Menge von Lokalpolitik-Parametern für jeden der Server (21-27), bis der Vergleichsschritt eine Übereinstimmung ergibt; und
- d) einer Einrichtung zum Speichern der Menge von Lokalpolitik-Parametern und der Mengen

von Diensteigenschaften als die Lokalpolitik jedes Servers für jeden der Server, wenn eine Übereinstimmung festgestellt wird;

eine Einrichtung, die durch den Makler (30, 31) auf der Grundlage der Netzpolitik und der Kapazität für verfügbare Betriebsmittel eines der Server den einen der Server für einen der Clients vorschlägt.

Revendications

1. Procédé mis en oeuvre par ordinateur pour affecter une multiplicité de serveurs (21 à 27) dans un réseau (5), chaque serveur ayant une capacité en ressources disponible, à une multiplicité de clients (11 à 19) pour fournir une multiplicité de services (A_1 à A_n) auxdits clients, le procédé comportant les étapes consistant:

a) à élaborer une politique de réseau pour ladite multiplicité de serveurs (21 à 27) en recueillant une politique locale pour chacun desdits serveurs, la politique locale pour chaque serveur étant basée sur une capacité globale de chaque serveur à fournir chaque service de la multiplicité de services (A_1 à A_n) auxdits clients (11 à 19);

b) à recevoir des demandes de clients pour lesdits services dans au moins un courtier (30, 31);
c) à faire une proposition de serveur, à partir du courtier (30, 31), à l'un desdits clients (11 à 19) faisant une demande de l'un desdits services (A_1 à A_n), la proposition de serveur identifiant un serveur proposé parmi lesdits serveurs sur la base de la politique de réseau et des capacités en ressources disponibles desdits serveurs, ledit serveur proposé parmi lesdits serveurs ayant la capacité en ressources disponible nécessaire à la fourniture de l'un desdits services;

d) à faire la demande de l'un desdits services (A_1 à A_n) de la part de l'un desdits clients (11 à 19) audit serveur proposé parmi lesdits serveurs (21 à 27) en réponse à la réception par ledit client parmi lesdits clients de la proposition de serveur de l'étape (c) de la part du courtier (30, 31);

e) à faire fonctionner le serveur proposé parmi lesdits serveurs conformément à sa politique locale respective pour rejeter la demande dudit service parmi desdits services (A_1 à A_n) de l'étape (d) lorsque la demande dépasse une limite de politique locale pour la politique locale respective du serveur proposé parmi lesdits serveurs; et

f) à faire fonctionner le serveur proposé parmi

lesdits serveurs (21 à 27) conformément à sa politique locale respective pour accepter la demande dudit service parmi lesdits services (A_1 à A_n) de l'étape (d) lorsque la demande ne dépasse pas la limite de politique locale pour la politique locale respective du serveur proposé parmi lesdits serveurs.

2. Procédé selon la revendication 1, dans lequel l'étape (c) consiste, en outre:

(a) à créer une liste de services (71), à l'intérieur dudit courtier (30, 31), des services disponibles à partir de ladite politique de réseau;

b) à créer une liste de serveurs respectifs (74) pour chaque service de la liste de services, chaque liste de serveurs contenant une entrée de serveur pour chaque serveur ayant la capacité globale nécessaire à la fourniture d'un service respectif; et

c) à contrôler, pour chaque service de la liste de services (71), un sous-ensemble desdites entrées de serveur dans ladite liste de serveurs pour proposer lesdits serveurs représentés dans ledit sous-ensemble auxdits clients.

3. Procédé selon la revendication 2, dans lequel ladite étape de contrôle dudit sous-ensemble comprend, en outre, les étapes consistant:

a) à établir un chemin de données (31 à 34), le chemin de données assurant la connectivité entre chaque serveur représenté dans le sous-ensemble et le courtier, le chemin de données transmettant des informations concernant la capacité en ressources disponible de chaque serveur représenté dans le sous-ensemble audit courtier; et

b) à déterminer si un premier serveur parmi les serveurs représentés dans le sous-ensemble a un compartiment client disponible (105 à 108) pour fournir le service demandé par lesdits clients sur la base de la capacité en ressources disponible dudit premier serveur.

4. Procédé selon la revendication 3, dans lequel ladite étape consistant à faire une proposition est exécutée lorsque l'un desdits serveurs faisant l'objet de contrôles a un compartiment client disponible (105 à 108) pour fournir ledit service.

5. Procédé selon la revendication 4, comportant en outre les étapes consistant:

a) à déconnecter le chemin de données (31 à 34) dudit premier serveur lorsque ledit chemin de données n'indique pas la présence de compartiments client disponibles (105 à 108); et

- b) à établir subséquemment avec le courtier (30, 31) un chemin de données partant d'un autre serveur de ladite multiplicité de serveurs de ladite liste de serveurs (74) pouvant prendre en charge ledit service et ayant des compartiments client disponibles.
6. Procédé selon la revendication 5, comportant en outre les étapes consistant:
- a) à élaborer une valeur de pondération d'analyse (73) pour chacun des serveurs de ladite multiplicité de serveurs (21 à 27) pouvant prendre en charge chaque service sur la base de ladite politique de réseau;
- b) à créer une entrée de pondération d'analyse contenant la valeur de pondération d'analyse pour chaque entrée de serveur de ladite liste de serveurs (74);
- c) à réduire la valeur de pondération d'analyse (73) pour un serveur sélectionné parmi lesdits serveurs représentés dans le sous-ensemble après que ledit serveur sélectionné a été proposé audit client; et
- d) à retirer ledit serveur sélectionné du sous-ensemble lorsque ladite valeur de pondération d'analyse est de zéro.
7. Procédé selon la revendication 1, dans lequel ladite étape d'élaboration d'une politique de réseau consiste:
- a) à déterminer un ensemble de caractéristiques de service pour chacun desdits services;
- b) à déterminer une capacité en services pour l'un desdits serveurs en fonction desdits ensembles de caractéristiques de service et d'un ensemble de paramètres de serveur dudit serveur parmi lesdits serveurs;
- c) à comparer la capacité en services à une performance souhaitée pour ledit serveur afin de déterminer s'il y a correspondance;
- d) à faire varier les ensembles de caractéristiques de service et l'ensemble de paramètres de serveur jusqu'à ce que ladite étape de comparaison produise une correspondance;
- e) à emmagasiner ledit ensemble de paramètres de serveur et ledit ensemble de caractéristiques de service en tant que politique locale pour ledit serveur lorsqu'une correspondance est déterminée;
- f) à répéter les étapes (a) à (e) pour chaque serveur de ladite multiplicité de serveurs; et
- g) à recueillir lesdites politiques locales pour obtenir une politique de réseau.
8. Procédé selon la revendication 1, dans lequel ledit réseau informatique (5) utilise de multiples courtiers (30, 31).
9. Procédé selon la revendication 8, dans lequel lesdits multiples courtiers (30, 31) constituent des répliques exacts les uns des autres et fonctionnent indépendamment.
10. Dispositif pour affecter une multiplicité de serveurs (21 à 27), ayant chacun une capacité en ressources disponible, à une multiplicité de clients (11 à 19) pour fournir un service parmi une multiplicité de services (A_1 à A_n) auxdits clients, lesdits serveurs et lesdits clients étant disposés dans un réseau informatique (5), le dispositif proposant un serveur proposé parmi la multiplicité de serveurs (21 à 27) à un client demandeur parmi la multiplicité de clients (11 à 19), le client demandeur parmi la multiplicité de clients (11 à 19) faisant la demande de l'un des services de la multiplicité de services (A_1 à A_n) à un serveur proposé parmi la multiplicité de serveurs fonctionnant de manière à rejeter la demande de l'un des services de la multiplicité de services lorsque la demande dépasse une limite de politique locale pour une politique locale respective du serveur proposé parmi la multiplicité de serveurs (21 à 27), le serveur proposé parmi la multiplicité de serveurs fonctionnant de manière à accepter la demande dudit service parmi la multiplicité de services (A_1 à A_n) lorsque la demande ne dépasse pas la limite de politique locale pour la politique locale respective du serveur proposé parmi la multiplicité de serveurs, le dispositif comportant:
- a) au moins un courtier (30, 31) comprenant:
- (i) des moyens pour recevoir une demande de client pour ledit service; et
- (ii) des moyens pour proposer de la part du courtier (30, 31) l'un desdits serveurs à l'un desdits clients faisant une demande sur la base d'une capacité en ressources disponible, ledit serveur proposé ayant la capacité en ressources disponible nécessaire à la fourniture dudit service.
11. Dispositif selon la revendication 10, dans lequel lesdits moyens de proposition de serveur comportent, en outre:
- 1) des moyens pour créer une liste de services (71), à l'intérieur dudit courtier (30, 31), de services disponibles;
- 2) des moyens pour créer une liste de serveurs (74) contenant des entrées de serveur des serveurs disponibles pour prendre en charge chacun desdits services de ladite liste de services; et

- 3) des moyens pour contrôler un sous-ensemble desdites entrées de serveur de ladite liste de serveurs (74) pour proposer lesdits serveurs représentés dans ladite liste de serveurs auxdits clients (11 à 19). 5
12. Dispositif selon la revendication 10 ou 11, dans lequel lesdits moyens de contrôle comprennent, en outre: 10
- a) un chemin de données (31 à 34) établissant la connectivité entre chaque serveur d'un ensemble de serveurs pouvant prendre en charge l'un desdits services et ledit courtier (30, 31), le chemin de données indiquant des compartiments client disponibles (105 à 108) pour le serveur audit courtier; et 15
- b) des moyens pour déterminer si ledit premier serveur a un compartiment client disponible pour fournir le service demandé par les clients. 20
13. Dispositif selon la revendication 12, dans lequel lesdits moyens de proposition envoient un message audit client lorsque l'un desdits serveurs a un compartiment client disponible pour la fourniture dudit service. 25
14. Dispositif selon la revendication 13, comportant, en outre: 30
- a) une valeur de pondération d'analyse (73) élaborée pour chacun des serveurs de ladite multiplicité de serveurs (21 à 27);
- b) une entrée de pondération d'analyse contenant la valeur de pondération d'analyse emmagasinée dans ledit courtier (30, 31) pour chaque entrée de serveur de ladite liste de serveurs; 35
- c) des moyens pour réduire la valeur de pondération d'analyse (73) pour ledit premier serveur dans la fenêtre de pré-visualisation (72) après que ledit premier serveur a été proposé audit client; et 40
- d) des moyens pour retirer ledit premier serveur de la fenêtre de pré-visualisation (72) lorsque ladite valeur de pondération d'analyse est de zéro. 45
15. Dispositif selon la revendication 12, dans lequel le dispositif comporte, en outre: 50
- des moyens pour élaborer une politique de réseau, comprenant:
- a) des moyens pour déterminer une capacité en services pour chaque serveur (21 à 27) sur la base d'une multiplicité d'ensembles de caractéristiques de service et d'un ensemble de paramètres de politique locale; 55
- b) des moyens pour comparer, pour chacun desdits serveurs, la capacité en services à une performance souhaitée pour chaque serveur afin de déterminer s'il existe une correspondance;
- c) des moyens pour faire varier, pour chacun desdits serveurs (21 à 27), les ensembles de caractéristiques de service et l'ensemble de paramètres de politique locale jusqu'à ce que ladite étape de comparaison produise une correspondance; et
- d) des moyens pour emmagasiner, pour chacun desdits serveurs (21 à 27), ledit ensemble de paramètres de politique locale et lesdits ensembles de caractéristiques de service en tant que politique locale pour chaque serveur lors de la détermination d'une correspondance;
- des moyens pour proposer, de la part du courtier (30, 31), l'un desdits serveurs à l'un desdits clients sur la base de la politique de réseau et de la capacité en ressources disponible dudit serveur parmi lesdits serveurs.




Figure 1




Figure 2




Figure 2A

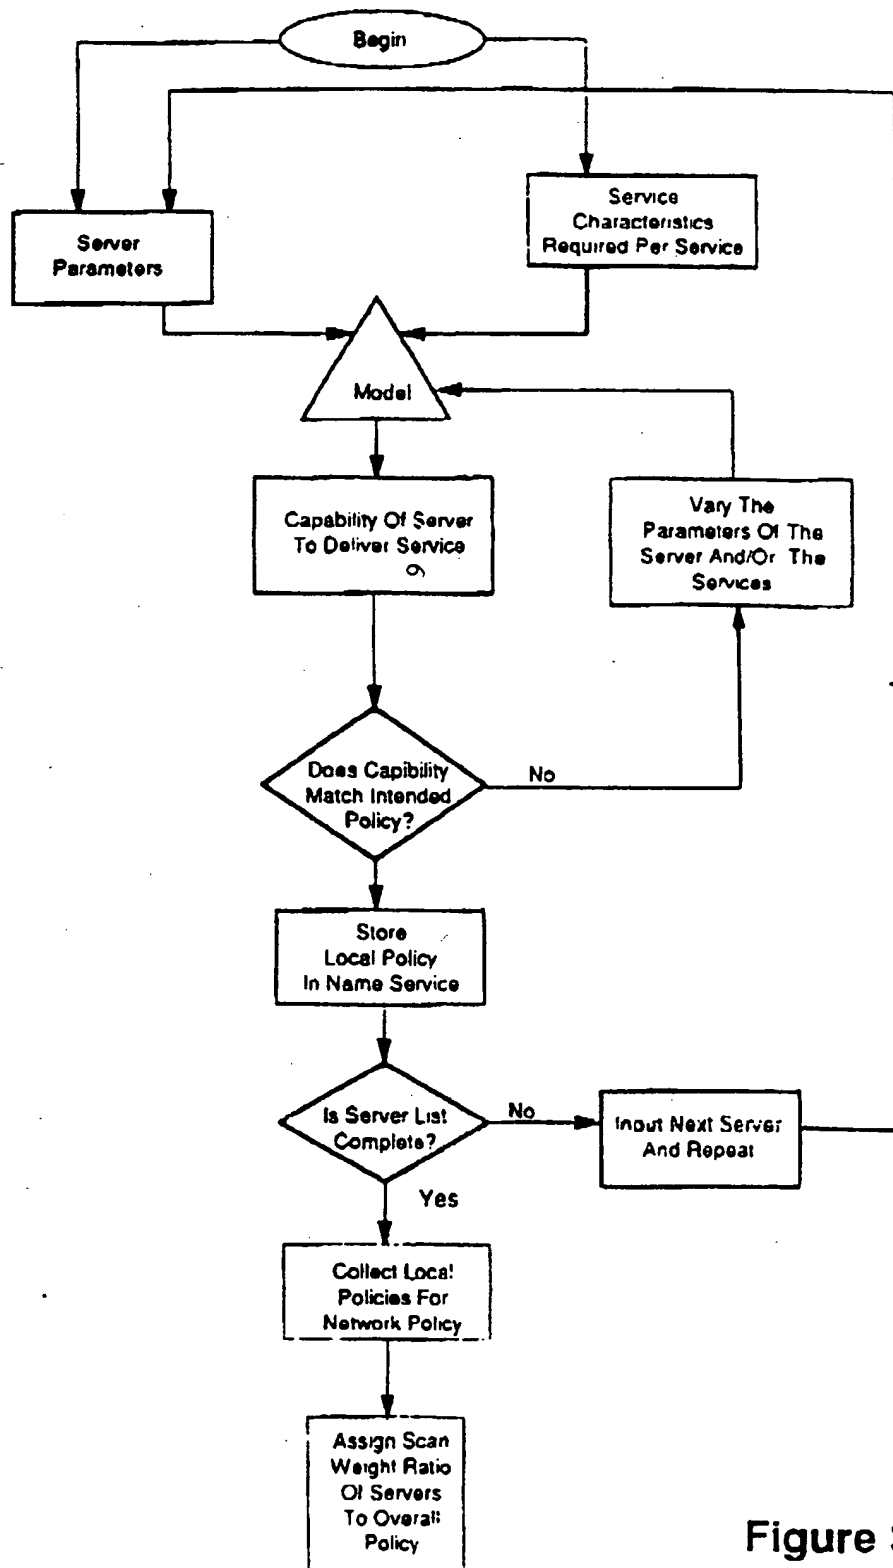


Figure 3




Figure 4

BEST AVAILABLE COPY




Figure 5




Figure 5A

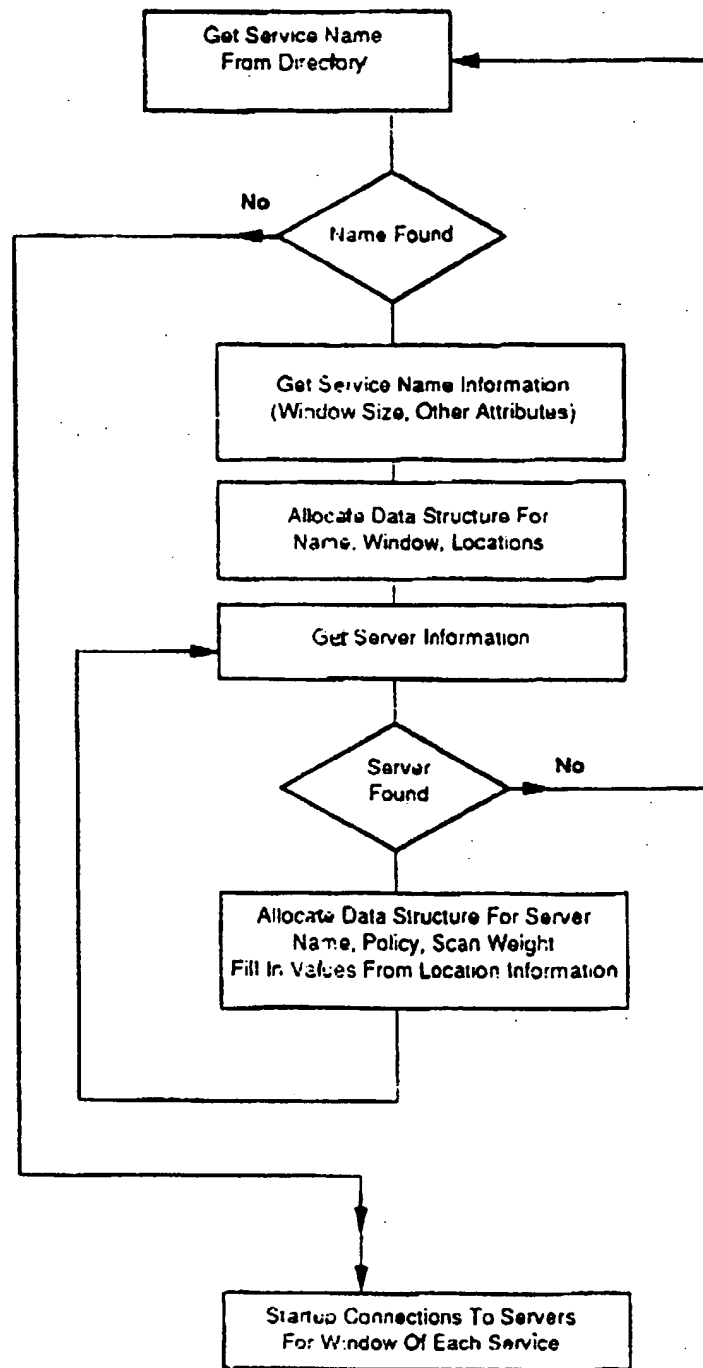


Figure 6

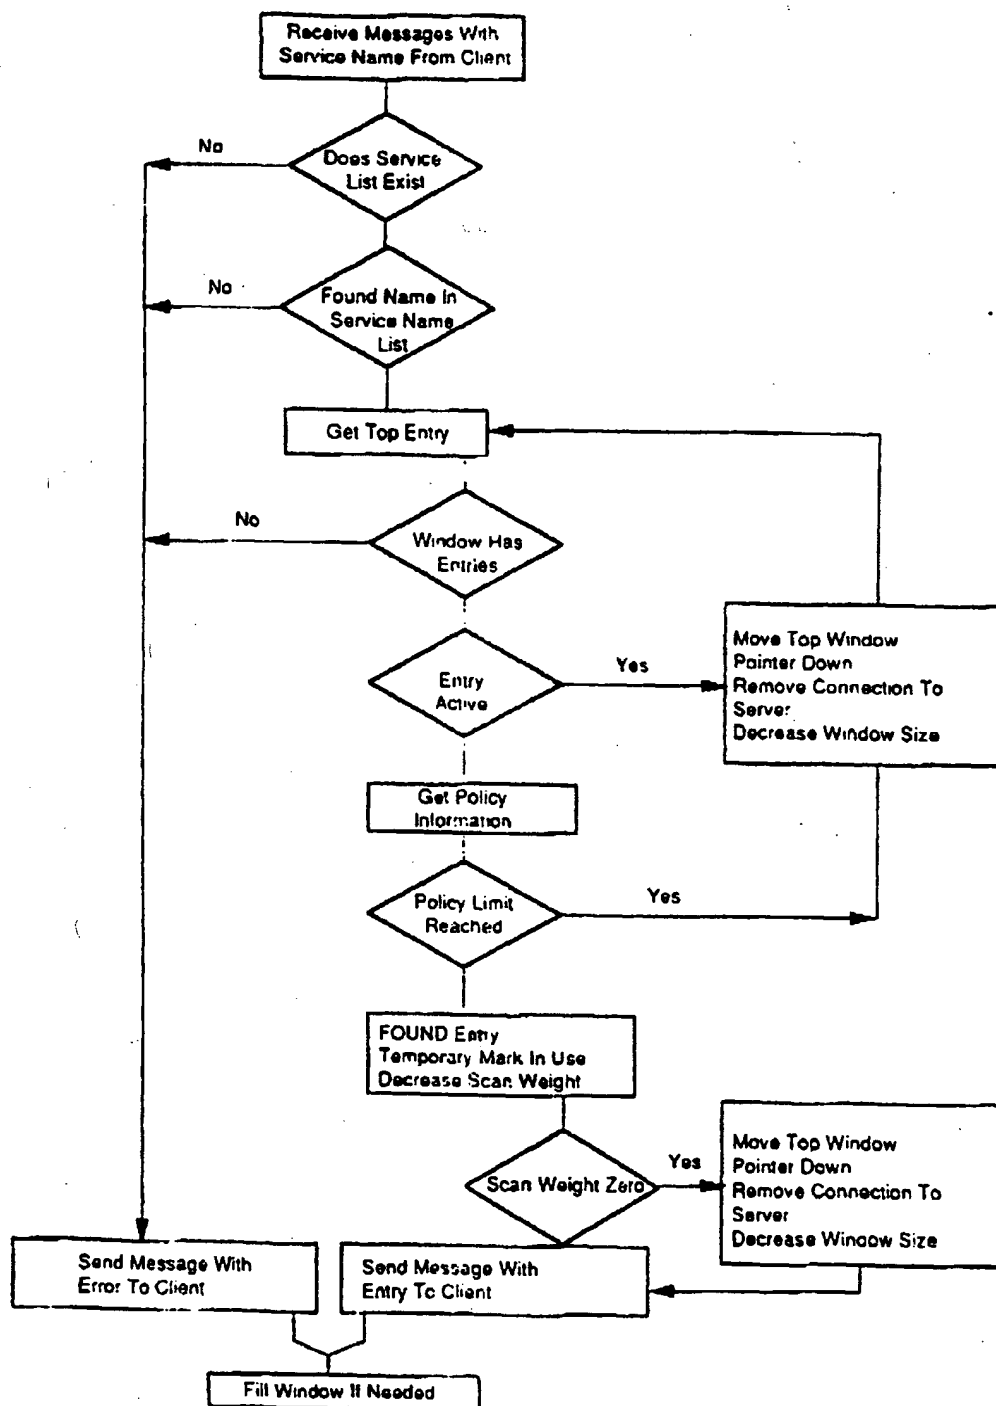


Figure 6A




Figure 7